

What is claimed is:

- 1 1. A method of selecting a heuristic class for data placement in a distributed
2 storage system comprising the steps of:
3 forming an integer program for each of a plurality of heuristic classes,
4 each of the integer programs comprising an objective of minimizing a
5 replication cost;
6 solving each of the integer programs which provide the replication cost
7 for each of the heuristic classes; and
8 selecting the heuristic class having a low replication cost.
- 1 2. A method of selecting a heuristic class for data placement in a distributed
2 storage system comprising the steps of:
3 forming a general integer program which models the data placement;
4 forming a specific integer program which models a heuristic class for
5 the data placement, the general and specific integer programs each
6 comprising an objective of minimizing a replication cost;
7 solving the general integer program which provides a general lower
8 bound for the replication cost;
9 solving the specific integer program which provides a specific lower
10 bound for the replication cost; and
11 selecting the heuristic class if a difference between the general lower
12 bound and the specific lower bound is within an allowable amount.
- 1 3. The method of claim 2 wherein inputs used in the steps of forming the general
2 and specific integer programs comprise a system configuration, a workload, and a
3 performance requirement.
- 1 4. The method of claim 3 wherein the performance requirement comprises a bi-
2 modal performance metric.
- 1 5. The method of claim 4 wherein the bi-modal performance metric comprises a
2 criterion and a ratio of successful attempts to total attempts.

- 1 6. The method of claim 3 wherein the performance requirement comprises a data
2 access latency.
- 1 7. The method of claim 3 wherein the performance requirement comprises a data
2 access bandwidth.
- 1 8. The method of claim 3 wherein the performance requirement comprises a data
2 update time.
- 1 9. The method of claim 3 wherein the performance requirement comprises an
2 average data access latency.
- 1 10. The method of claim 3 wherein the performance requirement comprises a data
2 availability requirement.
- 1 11. The method of claim 3 wherein the general integer program comprises general
2 constraints which model the data placement irrespective of the heuristic class for
3 the data placement.
- 1 12. The method of claim 11 wherein the general constraints comprise a
2 performance constraint which models the performance requirement.
- 1 13. The method of claim 11 wherein the specific integer program comprises the
2 general constraints and a specific constraint.
- 1 14. The method of claim 12 wherein the specific constraint comprises a storage
2 constraint.
- 1 15. The method of claim 12 wherein the specific constraint comprises a replica
2 constraint.
- 1 16. The method of claim 12 wherein the specific constraint comprises a routing
2 knowledge constraint and further wherein the routing knowledge constraint
3 models an extent to which a data storage node knows of replicas of data objects

4 stored on other data storage nodes.

1 17. The method of claim 12 wherein the specific constraint comprises an access
2 knowledge constraint and further wherein the access knowledge constraint models
3 an extent to which a data storage knows of access to replicas of data objects by
4 clients accessing other data storage nodes.

1 18. The method of claim 12 wherein the specific constraint comprises an activity
2 history constraint.

1 19. The method of claim 12 wherein the specific constraint comprises a reactive
2 placement constraint.

1 20. The method of claim 3 wherein the system configuration comprises a plurality
2 of data storage nodes coupled by a plurality of network links.

1 21. The method of claim 20 wherein the system configuration further comprises a
2 plurality of clients coupled to the data storage nodes.

1 22. The method of claim 21 wherein the workload comprises at least some of the
2 clients requesting data objects stored on the data storage nodes.

1 23. The method of claim 22 wherein the workload further comprises at least some
2 of the clients storing some of the data objects on the data storage nodes.

1 24. A method of selecting a heuristic class for data placement in a distributed
2 storage system comprising the steps of:
3 forming a general integer program which models the data placement;
4 forming a plurality of specific integer programs which model a
5 plurality of heuristic classes, the general and specific integer programs
6 each comprising an objective of minimizing a replication cost;
7 solving the general integer program which provides a lower bound for
8 the replication cost;
9 solving the specific integer programs which provides the replication

10 cost for each of the heuristic classes; and
11 selecting a particular heuristic class correlated to a low replication cost
12 if a difference between the lower bound and the low replication cost is
13 within an allowable amount.

1 25. A computer readable memory comprising computer code for implementing a
2 method of selecting a heuristic class for data placement in a distributed storage
3 system, the method of selecting the heuristic class comprising the steps of:
4 forming an integer program for each of a plurality of heuristic classes,
5 each of the integer programs comprising an objective of minimizing a
6 replication cost;
7 solving each of the integer programs which provide the replication cost
8 for each of the heuristic classes; and
9 selecting the heuristic class having a low replication cost.

1 26. A computer readable memory comprising computer code for implementing a
2 method of selecting a heuristic class for data placement in a distributed storage
3 system, the method of selecting the heuristic class comprising the steps of:
4 forming a general integer program which models the data placement;
5 forming a specific integer program which models a heuristic class for
6 the data placement, the general and specific integer programs each
7 comprising an objective of minimizing a replication cost;
8 solving the general integer program which provides a general lower
9 bound for the replication cost;
10 solving the specific integer program which provides a specific lower
11 bound for the replication cost; and
12 selecting the heuristic class if a difference between the general lower
13 bound and the specific lower bound is within an allowable amount.

1 27. A computer readable memory comprising computer code for implementing a
2 method of selecting a heuristic class for data placement in a distributed storage
3 system, the method of selecting the heuristic class comprising the steps of:
4 forming a general integer program which models the data placement;
5 forming a plurality of specific integer programs which model a

6 plurality of heuristic classes, the general and specific integer programs
7 each comprising an objective of minimizing a replication cost;
8 solving the general integer program which provides a lower bound for
9 the replication cost;
10 solving the specific integer programs which provides the replication
11 cost for each of the heuristic classes; and
12 selecting a particular heuristic class correlated to a low replication cost
13 if a difference between the lower bound and the low replication cost is
14 within an allowable amount.